arXiv:1702.01293v1 [cs.LG] 4 Feb 2017

Latent Hinge-Minimax Risk Minimization for Inference from a Small Number of Training Samples

Dolev Raviv
University of Haifa
dolev.raviv@gmail.com

Margarita Osadchy
University of Haifa
rita@cs.haifa.ac.il

Abstract

Deep Learning (DL) methods show very good performance when trained on large, balanced data sets. However, many practical problems involve imbalanced data sets, or/and classes with a small number of training samples. The performance of DL methods as well as more traditional classifiers drops significantly in such settings. Most of the existing solutions for imbalanced problems focus on customizing the data for training. A more principled solution is to use mixed Hinge-Minimax risk *[19] specifically designed to solve binary problems with imbalanced training sets. Here we propose a Latent Hinge Minimax (LHM) risk and a training algorithm that generalizes this paradigm to an ensemble of hyperplanes that can form arbitrary complex, piecewise linear boundaries. To extract good features, we combine LHM model with CNN via transfer learning. To solve multi-class problem we map pre-trained category-specific LHM classifiers to a multi-class neural network and adjust the weights with very fast tuning. LHM classifier enables the use of unlabeled data in its training and the mapping allows for multi-class inference, resulting in a classifier that performs better than alternatives when trained on a small number of training samples.*

1. Introduction

Many real binary classification problems involve imbalanced classes, for example object detection in vision and fraud detection in security. In such problems it is easy to collect background data, while data representing the target class is rare or hard (expensive) to obtain. The majority of existing powerful classifiers (e.g., SVM, Neural Networks, including deep ones) assume balanced training sets and when trained on imbalanced sets show degraded classification performance.

Deep Neural Networks have recently shown very impressive performance in large-scale multi-class problems [14, 25, 23, 24]. However, these models require very large number of **labeled** training samples and their performance drops rapidly when the training set size gets smaller. Note that the requirement of large labeled sets is expensive in terms of data collection and training time. In practice, many learning problems require rapid inference from small amounts of data.

The aim of this work is to develop powerful and fast classifiers that improve in both tasks: 1) training in imbalanced setting that involve a small number of positive training samples and a large number of negative data points; 2) multi-class problems with a small number of labeled samples.

We follow the paradigm introduced in [20] that combines hinge risk for the smaller class and minimax risk [16, 10] for the larger class to address imbalanced classification problems. The mixed risk was used to train linear and kernel hybrid classifiers in [20]. Unfortunately, while being well understood and fast, linear classifiers do not solve all machine learning problems. Kernel methods show good classification results for highly non-linear problems, but they suffer from long running time and are not scalable to large tasks. To address these issues, [19] derived a hinge-minimax risk and an efficient training algorithm for intersection of K positive halfspaces. Such an intersection forms a convex set which limits the applicability of the classifier.

In this work, we generalize the hinge-minimax risk for an ensemble of linear classifiers, that can form arbitrary, piece-wise linear boundaries. We propose a training algorithm that minimizes this risk by simultaneously discovering the convex components in the positive class and building K-hyperplane models to separate each component from the negative class. The learning is done by alternating between finding the best partition of the data into hidden components and updating the model over this partition. We call our novel classifier the *Latent Hinge Minimax* (LHM) classifier, as it discovers the latent structure in the data and employs the Hinge-Minimax paradigm.

We show that in imbalanced setting the proposed LHM classifier outperforms other combinations of hyperplanes, including Neural Network (NN) with an equivalent architecture (NN can be viewed as a combination of hyper-

planes). The robustness of LHM to imbalanced problems can be explained by the use of the minimax risk [16, 10], that serves as a regularizer in training (since it utilizes the statistics of the entire class, as opposed to learning from small batches of examples).

The LHM training procedure is designed for binary problems. To apply it in a multi-class setting, we build one-against-all classifiers for all classes and combine them in a single model by mapping class specific LHM models to a multi-class NN with a matching architecture (see Section 4.2). We then use the cross-entropy loss to adjust the weights in the resulting *LHM-NN* combination.

To solve classification problems with a small number of training examples, it was suggested (e.g.,[5, 7, 17]) to combine a pre-trained CNN (trained on a much larger training set for a related classification problem) for feature extraction, with a classifier for the target problem. Such an approach was also referred to as *transfer learning*. If the classifier is implemented as a neural network, it enables an end-to-end training, which usually improves the results. We show that using LHM-NN in the transfer learning settings has significant benefits compared to NN, in both classification accuracy and training efficiency. The improved accuracy stems from the ability of LHM model to learn from unlabeled data. The fast convergence of the LHM-NN (just a handful of epochs) is due to a very good initialization of the upper layers with class specific LHM classifiers. Note that class specific LHM models can be trained in parallel while a distributed training of fully connected layers in NN is far from being trivial. Moreover, adding a new class to LHM-NN is fast and easy: train a classifier for the new class, map it to the corresponding LHM-NN architecture and run a very fast fine-tuning. Similarly to [15], which considered the transfer learning for the $n+1$ category from a fully trained n-category classifier, we use only a handful of training samples for tuning it. In contrast to [15], we do not restrict the new classifier to belong to the span of the previously learned n classifiers. This allows us greater flexibility in adding a new, non-related class to the multi-class model.

The method proposed here is different from the one-shot learning approach [12, 22], which attempts to find a mapping between target and source examples and apply it to the examples or to the model. LHM classifier learns the target concept from its examples, leveraging from unlabeled data in modeling the background statistics.

2. Background

We first address the settings in which the positive labeled class is much smaller then the negative class. It was shown that hinge loss [26, 29, 1, 2, 11] is computationally appealing when there are fairly small number of training samples, thus it could be used to measure the positive class risk within imbalanced problem settings. Alternatively, the minimax risk [16, 10] upper bounds the distribution that generates the instances-labels examples in the world. This approach is computationally appealing when there are (infinitely) many training examples, since it only utilizes their statistical properties, such as mean and covariance. Consequently, it could be employed as the negative class risk.

In this work we derive a mixed risk and an efficient training algorithm for a more general ensemble of hyperplanes. Our approach builds upon the mixed risk for the intersection of K-hyperplanes [19] which is briefly summarized in section 2.1.

2.1. K-hyperplane Hinge-Minmax Classifier

Let $(x,y)\sim D$ be a joint distribution of samples $x \in \mathbb{R}^n$ and labels $y \in \{-1,1\}$. Let D_{neg} be a marginal distribution of samples over the negative labels, and μ and Σ be its mean and covariance respectively. For simplicity, unless stated otherwise, for a linear classifier w which predicts $y = sign(w^T x)$, we assume that $b = 0$ (or absorbed by w).

Let w_j, $j = 1,..,K$ denote K hyperplanes. Let W be a $K \times d$ matrix with w_j as its jth column. A K-hyperplane Hinge-Minmax classifier (KHHM) is an intersection of positive half-spaces defined by these K hyperplanes.

Let $X^+ \triangleq \{x \in X : y = 1\}$, and $X^- \triangleq \{x \in X : y = -1\}$ denote the positive and negative training sets correspondingly and let m^+ be the size of X^+ and m^- be the size of X^-. Let $\hat{\mu}$ and $\hat{\Sigma}$ be the mean and covariance of D_{neg}, estimated using X^-. The KHHM training algorithm in [19] minimizes the empirical risk:

$$L(W_{KHHM}) = L^{M,-1}_{X^-}(W) + L^{H,1}_{X^+}(W) \tag{1}$$

where $L^{M,-1}_{X^-}(W) = \sup_{z\sim Z(\hat{\mu},\hat{\Sigma})} \Pr(z \in Q)$ is the minimax risk over the negative labels inside the intersection $Q \triangleq \{x : W^T x \geq \vec{0}$ (the zero vector)$\}$. It was shown in [19] that

$$\sup_{z\sim Z(\mu,\Sigma)} \Pr(W^T z > \vec{0}) = \frac{1}{1+d^2}$$

with $d^2 = \mu^T \tilde{W}(\tilde{W}^T \Sigma \tilde{W})^{-1} \tilde{W}^T \mu$, where $\tilde{W}$ is a submatrix of W containing hyperplanes that intersect in a point closest to μ scaled by Σ^{-1}.

The hinge part of the risk in Eq. 1 is defined as $L^H_{X^+}(W) = \sum_{x\in X^+} \ell(W;x,1)$, where $\ell(W;x,y) = \sum_j \max\{0, 1 - y w_j^T x\}$ is the K-hyperplane hinge loss [19].

3. Latent Hinge-Minmax Classifier

To accommodate classes that form non-convex or disjoint sets, we propose a new model, Latent Hinge-Minmax (LHM) classifier, and a training scheme that simultaneously

discovers the convex components in the positive class and learns the K-hyperlane models separating each convex component from the negative class.

We define the LHM classifier as a union of intersections of positive half-spaces. We assume that each intersection is composed of K hyperplanes: $W^i = [w_1^i, ..., w_K^i]$ and there are C components in the union. Let $W_{LHM} \triangleq (W^1, \ldots, W^C)$ denote the LHM model. Equivalently, we can define the LHM classifier as

$$f_{\text{LHM}}(x, W_{LHM}) = sign(\max_{i\in\{1..C\}}\{\min_{j\in\{1..K\}} {w_j^i}^T x\}).$$

3.1. Latent Hinge-Minimax Risk

We extend the hinge-minimax risk in Eq. 1, to contain multiple latent components and a hidden assignment variable. Specifically, we define a latent variable $\varphi(x) = i, \;\; i \in \{0, \ldots, C\}$ for each sample $(x, y) \in D$. We set $\varphi(x) = 0$ for all samples with the negative label. Since the assignment of negative training samples is constant during the training, we reduce the set of latent values to $\{1, \ldots, C\}$ for the simplicity of notation.
We define the LHM risk function as follows:

$$L_D(W_{LHM}; \varphi) = L^M_{\mu,\Sigma}(W_{LHM}; \varphi) + L^H_D(W_{LHM}; \varphi), \quad (2)$$

where

$$L^M_{\mu,\Sigma}(W_{LHM}; \varphi) = \Pr_{z\sim Z(\mu,\Sigma)}(z \in \bigcup_{i\in\{1..C\}} Q^i) \quad (3)$$

is the minimax part of the LHM risk and

$$L^H_D(W_{LHM}; \varphi) = \quad (4)$$
$$\mathbb{E}_{(x,y)\in D}\left[\sum_{i=1}^{C} \ell(W^i; x, y)\mathbb{1}\left[\varphi(x) = i\right]\right]$$

is the hinge part, where

$$\ell(W; x, y) = \max_{j\in\{1..K\}}\{\max\{0, \alpha - y w_j^T x\}\}$$

is the modified K-hyperplane hinge loss. This change is required to accommodate comparison between the different norms of the hyperplanes.

3.2. Empirical Risk

Each sample with positive label encounters a loss only in a single latent component, specified by its latent variable $\varphi(x)$ as per Eq. 4. Thus,we define a single positive sample loss as follows,

$$L(W^{\varphi(x)}; x, 1, \varphi(x))$$
$$= \frac{1}{m^+_{\varphi(x)}} L^M_{X^-}(W^{\varphi(x)}) + \lambda L^H_{X^+}(W^{\varphi(x)}; x)$$

where $L^M_{X^-}(W^{\varphi(x)}) = \sup_{z\sim Z(\hat{\mu},\hat{\Sigma})} \Pr(z \in Q^i)$ is constant for all positive examples with the same assignment (the mean and covariance are estimated from X^-) and $L^H_{X^+}(W^{\varphi(x)}; x) = \ell(W^{\varphi(x)}; x, 1)$.

Let $X^i \triangleq \{x \in X^+ : \varphi(x) = i\}$ define a subset of X^+. The empirical risk of a latent component i aggregates the sample loss over all samples in X^i:

$$L(W^i) = \sum_{x\in X^i}[\frac{1}{m^+_i} L^M_{X^-}(W^i) + \lambda L^H_{X^+}(W^i; x)] \quad (5)$$
$$= L^M_{X^-}(W^i) + \lambda \sum_{x\in X^i}\left[L^H_{X^+}(W^i; x)\right]$$

Finally, we define the empirical risk of the LHM model as the sum of empirical risks of all its latent components:

$$L(W_{LHM}; \varphi) = \sum_{i=1}^{C} L(W^i) \quad (6)$$
$$= \sum_{i=1}^{C}\left(L^M_{X^-}(W^i)\right) + \lambda \sum_{i=1}^{C}\left(\sum_{x\in X^i} L^H_{X^+}(W^i; x)\right)$$

By summing the risk of the components in Eq. 6, we upper bounded the expected minimax risk in Eq. 3 with $\sum_{i=1}^{C} \Pr_{z\sim Z(\mu,\Sigma)}(z \in Q^i)$.

3.3. LHM Training

The training aims to minimize the empirical risk in Eq. 6 over the parameters W_{LHM} and the hidden variables φ. Similarly to latent SVM [28], the complexity of the optimal assignment of samples to latent components is exponential. We propose an iterative algorithm, which reaches fast convergence and shows good results in practice. The algorithm iterates between two steps: First, given an assignment it produces a model W_{LHM}, second, it updates the latent variables $\varphi(x), \forall x \in X^+$ to better represent the latent structure of the data.

The **first** step updates the LHM model W^t_{LHM} in iteration t given the latent variables φ from iteration $t-1$. Namely, for each hidden component $i = 1, ..., C$, we find the hyperplanes W^i separating the training samples in X^i from D_{neg} by minimizing the empirical risk in Eq. 5. This risk is minimized by the training algorithm proposed in [19].

The **second** step updates the latent variable assignment, given the current W^t_{LHM}. For each positive sample, it finds the best component w.r.t. the risk in Eq 6. Specifically, the hinge risk for x is simply $\ell(W^i, x, 1)$. The minimax part of the assignment function for $x \notin Q^i$ should consider the probability that this point adds when it is included in the component i (as shown in Figure 1, left). For $x \in Q^i$, the minimax part should consider the amount of probability released when the component shrinks as a result of change in

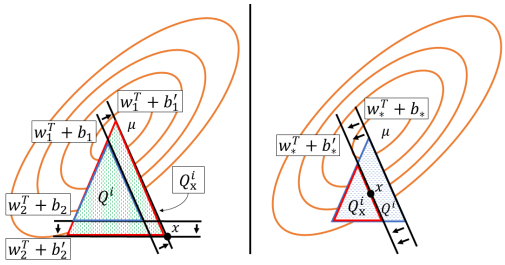

Figure 1. The orange elliptic circles represent the negative distribution $Z(\hat{\mu}, \hat{\Sigma})$, the red triangles corresponds to Q_x^i and the blue ones to Q^i. **Left**: w_1^i, w_2^i are moved to pass through x, causing the probability Q^i to increase. **Right**:w_*^i is moved to pass through x, causing the probability Q^i to decrease.

the assignment of x (as shown in Figure 1, right). The optimal assignment should take both cases into consideration for all components. We define the assignment as follows,

$$\varphi(x) = \tag{7}$$
$$\underset{i\in\{1..C\}}{\operatorname{argmin}} \left[\Pr_{z\sim Z(\hat{\mu},\hat{\Sigma})}(z \in Q_x^i) + \lambda \ell(W^i; x, 1) \right]$$

where $Q_x^i \triangleq \{x : {W_x^i}^T x \geq \vec{0}\}$ and

$$W_x^i \triangleq \begin{cases} W^{def} & \text{if } x \in Q \\ W^{inf} & \text{if } x \notin Q \end{cases}$$

W^{def} is a deflated model derived from W^i by parallel translation of the hyperplane closest to x such that $w_*^T x + b_* = 0$. W^{inf} is an inflated model derived from W^i by parallel translation of the hyperplanes for which $w_k^T x + b_k < 0$, until they intersect in x, namely, $w_k^T x + b_k' = 0$. The rest of the hyperplanes remain unchanged. The full training algorithm is summarized in Algorithm 1.

Algorithm 1 LHM Training. *KHHM-train* refers to the training of intersection of hyperplanes from [19]. T is the threshold on the empirical risk change.

Input: C, K, X^+, X^-, T
Output: W_{LHM}, φ

1: $t \leftarrow 1$
2: $L(W_{LHM}^{t=0}; \varphi^{t=0}) \leftarrow \infty$
3: $\varphi^t \leftarrow Init(X^+, C)$ $\{initial \quad assignment\}$
4: **do**
5: **for all** $i = 1, ..., C$ **do**$\{Model \quad Step\}$
6: $W^{i,t}$=*KHHM-training*(X^-,X^i)
7: **for all** $x \in X^+$ **do**$\{Assignment \quad Step\}$
8: $\varphi^{t+1}(x)$ as defined in Eq. 7
9: $t \leftarrow t + 1$
10: **while** $L(W_{LHM}^t; \varphi^t) - L(W_{LHM}^{t-1}; \varphi^{t-1}) \geq T$

Lemma 1. *Algorithm 1 minimizes the empirical risk* $L(W_{LHM}; \varphi)$.

Proof. Since LHM risk is a sum of risks over the latent components (Eq. 6), it is minimized by minimizing the empirical risk of each component. In step (5) of the Algorithm 1, we train $W^{i,t}$ model for each latent component $i = 1, ..., C$ using the iterative algorithm from [19] (the convergence of which was shown in [19]). It is easy to see that $L(W^i) = L(W_{KHHM})$, thus step (5) of the Algorithm 1 minimizes the component's risk in Eq. 5.

It is now left to show that the assignment φ^t in iteration t, will cause the reduction in the empirical risk in iteration $t + 1$. Since the empirical risk is aggregated over positive samples, it is enough to prove the claim for a single sample. We consider two cases:

1. The assignment of sample x does not change, formally $\varphi^t(x) = \varphi^{t+1}(x)$. In this case $L(W_{LHM}^{t+1}; \varphi^{t+1}(x))$ will only be affected by the $W^{i,t+1}$ training, thus

$$L\left(W_{LHM}^t; \varphi^t(x)\right) \geq L\left(W_{LHM}^{t+1}; \varphi^{t+1}(x)\right)$$

2. The assignment of sample x is changed. Formally in iteration t: $\varphi^t(x) = i$ and in interation $t + 1$ exists $j \neq i$, such that

$$\varphi^{t+1}(x) = j = \underset{k\in\{1..C\}}{\operatorname{argmin}} L_{X^-}^M\left(W_x^{k,t}\right) + \lambda L_{X^+}^H\left(W^{k,t}; x\right).$$

Since $x \in Q^i$, reassigning it to a different component will cause the $\Pr_{z\sim Z(\hat{\mu},\hat{\Sigma})}(z \in Q_x^i)$ to decrease (or stay the same), thus

$$L_{X^-}^M(W_x^{i,t}) - L_{X^-}^M(W^{i,t}) \leq 0.$$

Hence, the sample loss in component i is larger than the sample loss in the deflated component:

$$L\left(W^{i,t}; x\right) \geq L_{X^-}^M\left(W_x^{i,t}\right) + \lambda L_{X^+}^H\left(W^{i,t}; x\right). \tag{8}$$

At the same time, j is the optimal assignment, thus

$$L_{X^-}^M\left(W_x^{i,t}\right) + \lambda L_{X^+}^H\left(W^{i,t}; x\right) \geq \tag{9}$$
$$L_{X^-}^M\left(W_x^{j,t}\right) + \lambda L_{X^+}^H\left(W^{j,t}; x\right).$$

Since $W_x^{j,t}$ is a naive inflation of $W^{j,t}$ to include x, the solution $W^{j,t+1}$, provided by KHHM training, would have lower (or same) empirical risk, thus

$$L_{X^-}^M\left(W_x^{j,t}\right) \geq L_{X^-}^M(W^{j,t+1}). \tag{10}$$

In iteration $t + 1$, x is included in X^j for training the j'th latent component, consequently

$$L_{X^+}^H\left(W^{j,t}; x\right) \geq L_{X^+}^H(W^{j,t+1}; x). \tag{11}$$

(as we assume that $x \in X^j$ leads to $x \in Q^{j,t+1}$). Finally, by combining the inequalities in Eq. 8–11, we obtain:

$$L\left(W^{i,t}; x\right) \geq L\left(W^{j,t+1}; x\right).$$

□

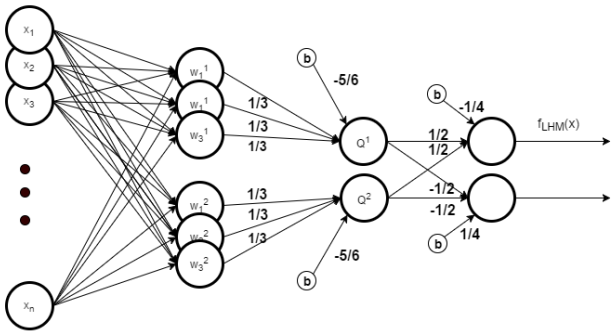

Figure 2. An example of NN equivalent to LHM.

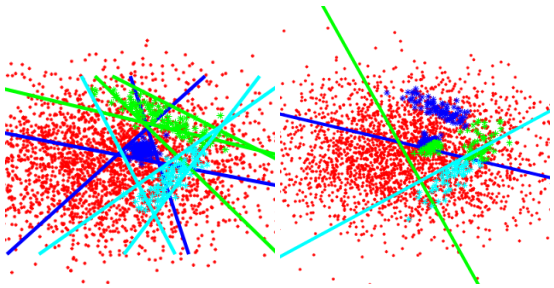

Figure 3. A qualitative comparison of the latent hinge minimax classifier (on the left) to the union of LDA classifiers (on the right).

4. Mapping LHM Classifier to a Neural Network

We propose to map LHM Classifier to a Neural Network. This enables 1) end-to-end training of the CNN features and LHM classifier for imbalanced problems and 2) LHM generalization to multi-class that enables using a smaller number of labeled training samples than NN.

4.1. Binary NN

A union of the intersections of positive half spaces can be implemented by a NN with three hidden layers. The first fully connected hidden layer has $K \times H$ neurons, where K is the number of hyperplanes in an intersection and H is the number of components. The second hidden layer has H nodes, connected only to the neurons associated with hyperplanes forming the corresponding intersection. The weights on these connections and the biases are fixed and mimic **AND** operation, namely, all weights of this layer are equal to $1/K$ and the biases are equal to $-1 + 1/(2K)$. The last hidden layer has two neurons, which are fully connected to the previous layer with the fixed weights and biases that mimic **OR** operation, namely, the fist neuron has weights equal to $1/H$ and the bias of $-1/(2H)$. The second output has weights equal to $-1/H$ and the bias of $1/(2H)$. The network has two outputs. An example of such network for $H = 2$ and $K = 3$ is depicted in Figure 2.

4.2. Multi-Class NN

For a multi-class setting, we suggest to train LHM model for each class using an additional unlabeled data for estimating the statistics of the negative class. We then map these models to a multi-class NN with the following architecture. The first hidden layer is a fully connected layer with $H \times K$ neurons per class, $H \times K \times C$ neurons in total, where C is the number of classes. These are equivalent to $H \times K \times C$ hyperplanes in the LHM model. For each hidden component, all hyperplanes in the intersection are connected to their corresponding node in the **AND** layer (as detailed in Section 4.1). The **AND** layer comprises $H \times C$ neurons. The next layer is a fully connected layer, comprising C nodes. The weights on the connections to the H components of the corresponding class are initialized with 1's, and the weights on the remaining connections are initialized with very small values from a Gaussian distribution. The network has C outputs and is trained using the cross-entropy loss.

To provide an end-to-end training, one can consider stacking the feature extraction layers of CNN (up to fully connected layers) with one of the above networks.

5. Experiments

We start by an illustrative example in 2D (Section 5.1) that shows the ability of the LHM classifier to discover the hidden components in the positive class and to separate each of them from the negative class using a K-hyperplane model.

Next, we compare LHM model to alternative ensembles of hyperplanes on the PASCAL-VOC 2007 dataset [6] (Section 5.2), and show its advantage over those methods and its robustness to the choice of the number of latent components. In these experiments we use simple HOG features and shallow architecture.

Finally, we show (Section 5.3) that LHM classifier can be combined with CNN via transfer learning. We address two settings: 1) binary problems with imbalanced sets, 2) multi-class tasks with a small number of labeled training samples. In both cases, LHM-based models show significantly better performance than NNs. The experiments are performed on images from cifar-10 and cifar-100 [13] and using LeNet CNN for features extraction.

5.1. Synthetic Data

A simpler alternative to the LHM model is a two-step algorithm which first finds the structure of the target class by applying some kind of unsupervised learning (e.g, k-means clustering) and then builds a model for each component. Such a simple approach was employed in [8] with LDA [9] classifier trained per cluster. Unless the clusters are very small (as in exemplar-based approach, which is time con-

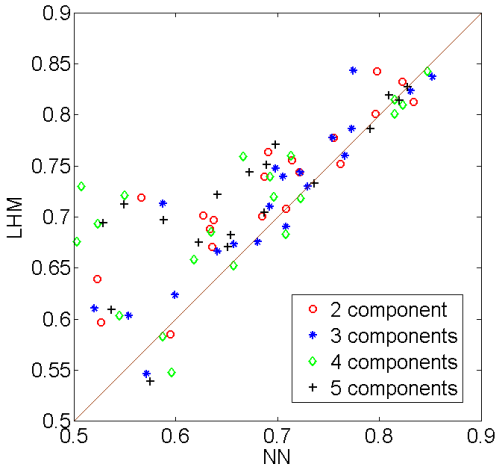

Figure 5. Comparison of the LHM classifier to the equivalent NN for a varying number of assumed hidden components (from 2 to 5) on PASCAL VOC 2007. The points above the diagonal line show the advantage of LHM classifier.

suming [18]), it relies heavily on the results of the clustering. If an initial clustering is incorrect (as in Figure 3, right), LDA (or any other convex classifier) cannot separate the resulting components from the background without including many false positives. The LHM training finds the underlying structure of the data and the model iteratively, improving both (Figure 3, left). Furthermore, LHM is quite robust to the initial assignment. Figure 4 shows a few iterations and the corresponding loss convergence when the initial assignment of the positive samples to components is chosen at random. Note the LHM training discovers the underlying structure in a 3-4 iterations.

5.2. Ensembles of Hyperplanes

Next, we compared the LHM classifier to alternative ensembles of linear classifiers on PASCAL VOC 2007 dataset [6] using Dalal-Triggs variant of the HOG features [3] with a fixed number of cells.

LHM model: We set the number of hyperplanes in each component to 2 and varied the number of components from 2 to 5. An initial assignment to the components was done using k-means with the Euclidian distance.

LDA Union (as a baseline model): We applied k-means clustering on whitened features to find the partition. We then learned an LDA classifier for each cluster in that partition. We varied the number of clusters from 2 to 5.

NN with an architecture equivalent to LHM: We used the model described in Section 4.1 with $K = 2$ and $H = 2, .., 5$, but the weights were initialized at random.

KHHM model [19]: This is essentially an LHM model with a single component, thus it is theoretically inferior to LHM. However, we ran this experiment to test the benefits of modeling the hidden structure of the positive class. We varied the number of hyperplanes from 2 to 5.

LHM	Union of LDAs	NN	KHHM
71.48%	65.17%	67.19%	69.45%

Table 1. The table reports (1-EER)*100 averaged over 20 classes and different hidden partitions (except for KHHM) on PASCAL VOC-2007 classification task using 80-dimensional HOG features.

All ensembles were trained in one-against-all manner. Similarly to [8, 21], we learned the background mean and covariance using bounding boxes from all classes and used them to represent the negative class in LDA union, KHHM, and LHM training.

We tested all ensemble classifiers on all windows from the test set. Table 1 summarizes the results ($(1-EER)\cdot 100$)) for all tested ensembles averaged over classes and different parameters. It shows that LHM model outperforms all other classifiers. Figure 5 compares LHM to NN on 20 categories (as one-against-all binary classifiers) for varying number of hidden components. The plot shows that LHM outperforms NN independently of the number of components.

5.3. Deep Architecture

Next, we tested the LHM classifier on top of the pre-trained CNN feature extraction in imbalanced binary problems and in multi-class tasks with a small number of labeled examples. We explored the following transfer learning settings. The first setting refers to the **best case** scenario in which the source and the target classification tasks operate on the *same* set of features but differ in the classification problem. The second setting refers to the **worst case** scenario for the transfer learning where the source and the target classification problems *share very little similarity*. The "worst case" scenario is very common in practice, as many classification tasks do not have a large, comprehensive training set (such as ImageNet [4] in object recognition) to be used in transfer learning. No good solution currently exists for such problems.

We used the CIFAR-10, composed of 10 categories (airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck) as the source problem. Specifically, we trained the LeNet model implemented in MatConvNet [27] on CIFAR-10. Then we removed the last fully-connected layer and the soft-max and used this trimmed network as a feature extractor which converts images to a 64-dimensional feature vectors.

For the best case transfer learning, we defined a new set of classes by coupling i and $i + 5$ indexes of CIFAR-10 classes. CNN trained on CIFAR-10 maps individual classes to linearly separable sub-spaces, thus using pairs of classes as a target classification problem makes it non-linear. Consequently, we get a new classification problem over the same space of features.

For the worst case transfer learning, we picked a subset

of 5 classes (train, bottle, cattle, forest, and sweet peppers) from the CIFAR-100, which do not overlap (in their visual appearance) with the CIFAR-10 categories, to be the target classification task. CIFAR-10 data set is not rich enough to enable learning of features that can be used for an arbitrary category, thus we believe that such setting is especially difficult.

We tested the LHM binary and multi-class classifiers in the best and the worst case transfer learning scenarios and compared their performance to two baselines. One is an NN with a single fully connected layer and the cross-entropy loss (NN linear) and the other is the NN with the architecture matching the LHM model (NN matching). We repeated each experiment 50 times over different random subsets of training samples and random initialization of NN and averaged the results.

5.3.1 Binary Imbalanced Problem

The "Best Case" Transfer Learning: We trained binary classifiers for pairs of classes from CIFAR-10 using imbalanced training sets, in which the negative class included all samples from all other classes (40,000 examples) and the positive class included a varying number of samples (140, 300, 600, 1400, 2000, 5000-all). This resulted in imbalance ratios from 1:256 to 1:4.

LHM model was trained with 2 hidden components and 3 hyperplanes per component. The matching NN mimicked the configuration of LHM model, but the weights were allowed to change in training. Figure 6-left shows the 1-EER (averaged over 5 classification problem) of the LHM classifier and the two NN baselines as a function of the positive training sample size.

The "Worst Case" Transfer Learning: Since the number of samples per class in CIFAR-100 is significantly smaller, this experiment tests the robustness to imbalanced training data and to a small number of examples. We varied the size of the positive training set between 20, 50, 100, 250, 500(all) samples and we used all 2,000 samples of other classes as the negative training set. We compared the LHM model trained with 2 hidden components and 2 hyperplanes per component to NN baselines. Figure 6-right shows the 1-EER of the classifiers averaged over 5 classification problems as a function of the positive training set size.

5.3.2 Multi-Class Problem

The "Best Case" Transfer Learning: We mapped the LHM binary classifiers trained for 5 pairs of categories to a multi-class NN as described in Section 4.2. We fine-tuned the weights with a very fast training (just a handful of epochs, while training from scratch requires two orders of magnitude more training epochs). Figure 7-left shows the accuracy of the LHM models mapped to a multi-class NN (LHM-NN) with the two baseline NNs as a function of the size of the training set.

The "Worst Case" Transfer Learning: We mapped the LHM binary classifiers trained for the 5 categories from CIFAR-100 (using CIFAR-10 features) to a multi-class NN and fine-tuned the weights with a small number of epochs.

To test the complexity of the transfer learning problem we also trained a CNN (LeNet model implemented in MatConvNet [27]) on the target problem. We hoped that due to the small size of the target classification problem, 500 training examples per class would yield relatively good accuracy. Figure 7-right compares the accuracy of LHM-NN, two baseline NNs, and CNN (trained from scratch) as a function of the training sample size. It shows that CNN trained on the target problem is indeed the best as it succeeds to learn features specific for the task, but its accuracy drops very abruptly when the number of training samples becomes smaller. This suggests that when the number of training examples is small, using transfer learning even in a such difficult setting is a better solution than training a CNN from scratch.

The results in Figures 6 and 7 show that the NN models either heavily overfit when the number of training samples is small (NN matching) or they are not expressive enough when the number of training samples increases (NN linear). LHM classifiers are expressive enough to learn from a large set of examples and are more robust to overfitting when the number of examples is small.

6. Training Efficiency

Another advantage of LHM-NN is its training efficiency. A class-specific LHM model converges in 5-10 iterations. Its training time primarily depends on the number of positive samples and the dimension. The negative samples are used to estimate the mean and covariance of the background. The initial estimation (which involves a large number of samples) can be done only once and used for all classes. Since the probability of the negative class is evaluated inside the positive region using false positives [19], the number of which drops very fast, the estimation time of the mean and covariance during the training is negligible. Training of a binary classifier per class is independent of other classes, thus their training can be done in parallel. Finally, the fine-tuning of the multi-class network after mapping is very fast, due to the initialization of all layers (using supervized learning): feature extraction layers with pre-trained CNN and classifier's layers with LHM models.

The LHM-NN is also beneficial for the problems in which classes are dynamically added or removed from the classification task. Adding a class requires training a single binary classifier and fast fine-tuning; removing a class requires only fine-tuning.

7. Conclusions

We proposed a novel Latent Hinge-Minimax classifier for binary problems that discovers the hidden components in the positive class and separates them from the negative class with the intersections of positive half spaces. The main advantage of this classifier is its ability to incorporate unlabeled data in training. This results in a better robustness to imbalanced problems. We showed that for multi-class tasks, class-specific LHM models can be mapped to a multi-class NN with matching architecture requiring only a few iterations of fine-tuning. Finally, the proposed LHM architecture can be integrated with CNN features via transfer learning. The entire training procedure is very efficient. Our experiments showed that such classifiers are much more robust to the number of labeled training samples than the equivalent NNs.

We plan to incorporate multi-class loss into the Hinge-Minimax paradigm and design an efficient algorithm for minimizing this risk. We also plan to tain NNs using Hinge-Minimax like loss. It would be interesting to compare the results of this model with the binary-to-multiclass mapping scheme proposed here.

References

[1] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002. 2

[2] O. Bousquet, S. Boucheron, and G. Lugosi. Introduction to statistical learning theory. In *Advanced lectures on machine learning*, pages 169–207. Springer, 2004. 2

[3] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, pages 886–893, 2005. 6

[4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 6

[5] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 647–655, 2014. 2

[6] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 5, 6

[7] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition,CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 580–587, 2014. 2

[8] B. Hariharan, J. Malik, and D. Ramanan. Discriminative decorrelation for clustering and classification. In *Computer Vision - ECCV 2012 - 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part IV*, pages 459–472, 2012. 5, 6

[9] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer New York Inc., 2001. 5

[10] J. Honorio and T. Jaakkola. {Tight Bounds for the Expected Risk of Linear Classifiers and PAC-Bayes Finite-Sample Guarantees}. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 384–392, 2014. 1, 2

[11] S. M. Kakade, K. Sridharan, and A. Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pages 793–800, 2009. 2

[12] G. Koch, R. Zemel, and R. Salakhutdinov. Siamese neural networks for one-shot image recognition. In *In ICML Deep Learning workshop*, 2015. 2

[13] A. Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Master's thesis, 2009. 5

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1

[15] I. Kuzborskij, F. Orabona, and B. Caputo. From N to N+1: multiclass transfer incremental learning. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, pages 3358–3365, 2013. 2

[16] G. R. Lanckriet, L. E. Ghaoui, C. Bhattacharyya, and M. I. Jordan. A robust minimax approach to classification. *Journal of Machine Learning Research*, 3(Dec):555–582, 2002. 1, 2

[17] E. Littwin and L. Wolf. The multiverse loss for robust transfer learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2

[18] T. Malisiewicz, A. Gupta, and A. A. Efros. Ensemble of exemplar-svms for object detection and beyond. In *IEEE International Conference on Computer Vision, ICCV 2011, Barcelona, Spain, November 6-13, 2011*, pages 89–96, 2011. 5

[19] M. Osadchy, T. Hazan, and D. Keren. K-hyperplane hinge-minimax classifier. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1558–1566, 2015. 1, 2, 3, 4, 6, 7

[20] M. Osadchy, D. Keren, and B. Fadida-Specktor. Hybrid classifiers for object classification with a rich background. In *European Conference on Computer Vision*, pages 284–297. Springer, 2012. 1

[21] M. Osadchy, D. Keren, and D. Raviv. Recognition using hybrid classifiers. *IEEE transactions on pattern analysis and machine intelligence*, 38(4):759–771, 2016. 6

[22] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. P. Lillicrap. Meta-learning with memory-augmented neural networks. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 1842–1850, 2016. 2

[23] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 1

[24] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich.

Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1–9, 2015. 1

[25] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014. 1

[26] V. N. Vapnik. The nature of statistical learning theory, ser. statistics for engineering and information science. *New York: Springer*, 21:1003–1008, 2000. 2

[27] A. Vedaldi and K. Lenc. MatConvNet: Convolutional Neural Networks for MATLAB. In *Proc. of the 23rd Annual ACM Conference on Multimedia Conference, Brisbane*, pages 689–692, 2015. 6, 7

[28] C.-N. J. Yu and T. Joachims. Learning structural svms with latent variables. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pages 1169–1176, New York, NY, USA, 2009. ACM. 3

[29] T. Zhang. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research*, 2(Mar):527–550, 2002. 2

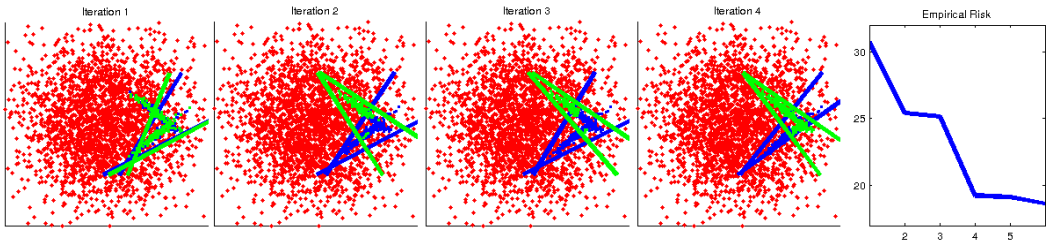

Figure 4. First four iterations of the LHM training on toy example and the corresponding loss convergence.

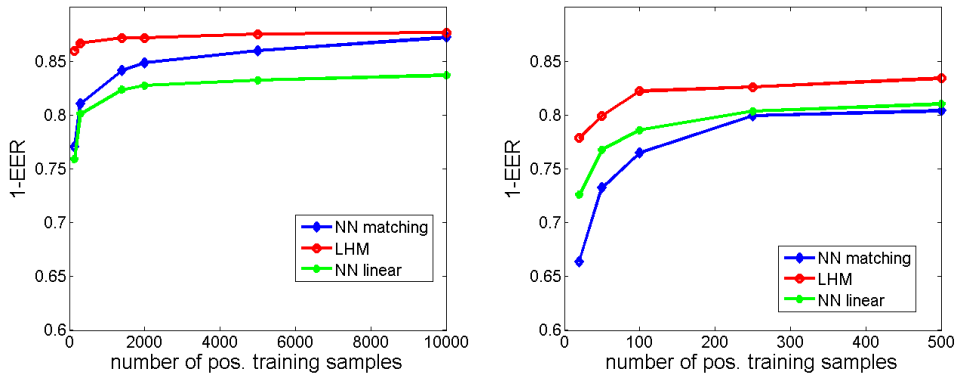

Figure 6. Binary imbalanced classification: left – the "best-case" transfer learning setting, right – the "worst-case" transfer learning setting.

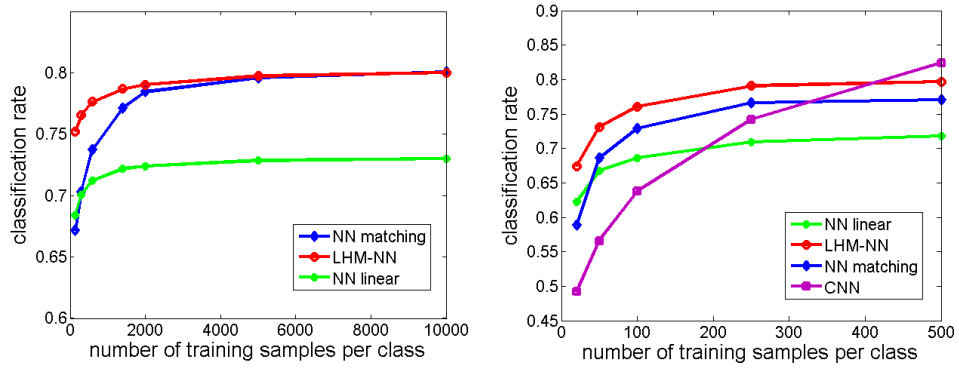

Figure 7. Multi-class classification: left – the "best-case" transfer learning setting, right – "worst-case" transfer learning setting.